

US-PAT-NO: 5604862

DOCUMENT-IDENTIFIER: US 5604862 A

TITLE: Continuously-snapshotted protection of computer files

----- KWIC -----

Abstract Text - ABTX (1):

An Integrity Server computer for economically protecting the data of a computer network's servers, and providing hot standby access to up-to-date copies of the data of a failed server. As the servers' files are created or modified, they are copied to the Integrity Server. The invention provides novel methods for managing the data stored on the Integrity Server, so that up-to-date snapshots of files of the protected file servers are stored on low-cost media such as tape, but without requiring that a system manager load large numbers of tapes.

TITLE - TI (1):

Continuously-snapshotted protection of computer files

Brief Summary Text - BSTX (3):

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office file or records, but otherwise reserves all copyright rights whatsoever.

Brief Summary Text - BSTX (6):

Known computer backup methods copy files from a computer disk to tape. In a full backup, all files of the disk are copied to tape, often requiring that all users be locked out until the process completes. In an "incremental backup," only those disk files that have changed since the previous backup, are copied to tape. If a file is corrupted, or the disk or its host computer fails, the last version of the file that was backed-up to tape can be restored by mounting the backup tape and copying the backup tape's copy over the corrupted disk copy or to a good disk.

Brief Summary Text - BSTX (11):

The invention provides methods and apparatus for protecting computer data against failure of the storage devices holding the data. The invention provides this data protection using hardware and storage media that is less expensive than the redundant disks required for disk mirroring, and protects against more types of data loss (for instance, user or program error) while providing more rapid access to more-recent "snapshots" of the protected files than is typical of tape backup copies.

Brief Summary Text - BSTX (12):

In general, in a first aspect, the invention features a method for managing copies of a protected set of files on a bounded number of volumes of sequential-access media. In the method, one of the sequential-access media volumes is chosen as the current volume. When the contents of one of the protected files is altered, the new current version is copied to the current

volume. When the current volume is full to a defined limit, a new volume is selected to be the current volume. The population of an active set (the minimum set of the most-recently-current of the volumes that together contain at least one version of each of the protected files) of the sequential-access volumes is maintained at or below the bounded number by periodically selecting a volume from the active set (typically the oldest) for compaction, and copying from the compaction volume to the current volume those versions of file versions stored on the compaction volume not having a more recent version stored on the active set. The copying and compacting steps continue while client nodes continue to alter the files of the servers.

Brief Summary Text - BSTX (13):

Preferred embodiments of the first aspect may feature the following. New versions of protected files, and versions reclaimed from compaction volumes, are copied to a direct access storage cache and queued for later writing to the current volume. When file versions are dequeued, the queue is reviewed for later versions of the dequeued file: only the latest version of the dequeued file is actually written to the active volume, and other versions in the queue are purged. Storage records are maintained to record the storage locations of file versions in the storage volumes so that the file versions can be accessed promptly. Recently-compacted volumes are maintained as a legacy set of volumes containing additional copies of current versions and non-current versions of files, and the storage records are maintained to track the contents of the legacy volumes. The volumes are cartridge tapes kept in an autoloader, and tape mounts/dismounts are automatically scheduled by software. A second set of volumes is also written concurrently; this set of volumes contains less-frequent snapshots than the active set, and a policy ensures that at least one version of each of the protected files is copied to the archival set within a bounded maximum interval.

Brief Summary Text - BSTX (14):

In a second aspect, the invention provides a method for protecting the data files of a computer, as the files are created and altered by an external process. In the method, recently-altered protected files are snapshotted to a storage cache. A new snapshot of a given file displaces any older snapshot of the same file from the storage cache. Later, non-displaced snapshotted versions are copied from the storage cache to removable mass storage media. This second copying phase proceeds at a lower rate, so that a significant proportion of the snapshotted versions of rapidly-changing files are displaced from the archive storage cache.

Brief Summary Text - BSTX (15):

Preferred embodiments may include the following features. The protected direct-access mass storage device includes the individual mass storage devices of file server nodes of a computer network. The content of the stored snapshots periodically verified against the protected files. This verification may use a technique that avoids copying contents of verified files over the network, or reading the removable media, by comparing a summary value of the content of a protected file with a summary value of the content of the stored snapshot. Stable protected files, those not recently altered, are periodically snapshotted, thereby to generate a media archive of all protected files, suitable for off-site storage. Periodically, the off-site media are selectively expired, leaving short sequences of consecutive media that, taken together, store at least one copy of every protected file. Various scheduling policies are available for the snapshotting: continuous scanning, at a specified time of day, in response to specific system events, or on demand.

Brief Summary Text - BSTX (16):

In a third aspect, the invention features a method in which files of a file system are traversed for snapshotting to removable storage media. A record is kept of the files currently held open for snapshotting. When an client process requests access to a file, the record is consulted to determine whether the file is currently open by the protection process. If the file is currently held open by the protection process, the client is blocked until the protection process releases the file. If the file is not currently held open by the protection process, or when the protection process completes the snapshotting, the file is opened in accord with the file open protocol of the protected computer.

Brief Summary Text - BSTX (17):

The invention has many advantages, including the following. A nearly-up-to-date copy of every file of the protected set is always available in the storage cache or the removable media. The snapshots can be used either to restore an image of a protected server if the server fails, or a user can get access to historical snapshots of files, for instance to compare the current version of a file to a version for a specified prior time. An ordinary user can, in seconds, access any file snapshot that was stored on an unavailable server node, or can request a restore of any version snapshot available to the Integrity Server.

Brief Summary Text - BSTX (18):

The active set can replace daily incremental backup tapes, to restore the current or recent versions of files whose contents are corrupted or whose disk fails. Note, however, that the data on the active set has been sampled at a much finer rate than the data of a daily backup. Thus, a restore recovers much more recent data than the typical restore from backup.

Detailed Description Text - DETX (4):

Referring to FIGS. 1, 2a, and 2b, the Integrity Server system operates in two main modes: protection mode and stand-in mode. When all file servers 102 under the protection of Integrity Server 100 are operational (FIGS. 1 and 2a), the system operates in protection mode: Integrity Server 100 receives up-to-date copies of the protected files of the servers 102. When any protected server 102 goes down (FIGS. 1 and 2b), the system operates in stand-in mode: Integrity Server 100 provides the services of the failed server 102, while still protecting the remaining protected servers 102. The software is divided into three main components: the agent NLM (NetWare Loadable Module) that runs on the server nodes 102, the Integrity Server NLM that runs on the Integrity Server 100 itself, and a Management Interface that runs on a network manager's console as a Windows 3.1 application.

Detailed Description Text - DETX (6):

After a client node 104 updates a file of a file server 102, producing a new version of the file, the agent process on that file server 102 copies the new version of the file to the Integrity Server's disk 120. As the file is copied, a history package 140 is enqueued at the tail of an active queue 142 in the Integrity Server's storage 130; this history package 140 holds the data required for the Integrity Server's bookkeeping, for instance telling the original server name and file pathname of the file, its timestamp, and where the Integrity Server's current version of the file is stored. History package 140 will be retained in one form or another, and in one location or another (for instance, in active queue 142, offsite queue 160, or the catalog--see

FIGS. 3a -3b) for as long as the file version itself is managed by Integrity Server 100.

Detailed Description Text - DETX (7):

When history package 140 reaches the head of active queue 142, the file version itself is copied from disk 120 to the current tape 150 in autoloader 110. History package 140 is dequeued to two places. History package 140 is enqueued to off-site queue 160 (discussed below), and is also stored as history package 312 in the protected files catalog, in a format that allows ready lookup given a ".backslash.backslash.server.backslash.file" pathname, to translate that file pathname into a tape and an address on that tape at which to find the associated file version.

Detailed Description Text - DETX (8):

As tape 150 approaches full, control software unloads current tape 150 from the autoloader read/write station, and loads a blank tape as the new current tape 150. The last few current tapes 151-153 (including the tape 150 recently removed, now known as tape 151) remain in the autoloader as the "active set" so that, if one of servers 102 fails, the data on active set 150-153 can be accessed as stand-in copies of the files of the failed server 102.

Detailed Description Text - DETX (9):

When a file version is written to active tape 150, its corresponding history package 140 is dequeued from active queue 142 and enqueued in off-site queue 160. When an off-site history package 162 reaches the head of off-site queue 160, the associated version of the file is copied from disk 120 to the current off-site tape 164, and the associated history package 312 is updated to reflect the storage of the data to offsite media in the protected file catalog. History package 312 could now be deleted from disk 120. When current off-site tape 164 is full, it is replaced with another blank tape, and the previous off-site tape is removed from the autoloader, typically for archival storage in a secure off-site archive, for disaster recovery, or recovery of file versions older than those available on the legacy tapes.

Detailed Description Text - DETX (10):

The size of the active tape set 150-153 is fixed, typically at three to four tapes in a six-tape autoloader. When a new current tape 150 is about to be loaded, and the oldest tape 153 in the set is about to be displaced from the set, the data on oldest tape 153 are compacted: any file versions on tape 153 that are up-to-date with the corresponding files on protected servers 102 are reclaimed to disk cache 120, from where the file will again be copied to the active and off-site tapes. Remaining file versions, those that have a more-recent version already on tapes 150-152 or on disk 120, are omitted from this reclamation. Once the data on tape 153 has been reclaimed to disk 120, tape 153 can be removed from the autoloader and stored as a legacy tape, typically either kept on-site for a few days or weeks before being considered blank and reused as a current active tape 150 or off-site tape 164, or retained for years as an archive. The data reclaimed from tape 153 are copied from disk 120 to now-current tape 150. The reclaimed data are then copied to tape 164 as previously described. This procedure not only maintains a compact number of active tapes, but also ensures that a complete set of data from servers 102 will appear in a short sequence of consecutive offsite tapes, without requiring recopying all of the data from the servers 102 or requiring access to the offsite tapes.

Detailed Description Text - DETX (11):

Referring to FIG. 2a, as noted earlier, as long as all servers 102 are functioning normally, all clients 104 simply read and write files using normal network protocols and requests, and agent processes on each of the servers 102 periodically copy all recently-modified files to Integrity Server 100. Integrity Server 100, at least in its role of protecting file servers 102, is essentially invisible to all clients 104.

Detailed Description Text - DETX (12):

Referring to FIG. 2b, after one of servers 202 fails, Integrity Server 100 enters stand-in mode (either automatically or on operator command). Integrity Server 100 immediately begins building a replica of the protected server's volume and directory structure, using the information stored on disk 120 and tapes 150-153, 164. Integrity Server 100 assumes the identity of failed server 202 during connect requests, intercepts network packets sent to failed server 202, and provides most of the services ordinarily provided by failed server 202. Clients 104 still request data from failed server 202 using unaltered protocols and requests. However, these requests are actually serviced by Integrity Server 100, using the replica of the failed server's file system. This stand-in service is almost instantaneous, with immediate access to recently-used files, and a few seconds' delay (sometimes one or two seconds, usually within a minute, depending on how near the tape data are to the read/write head) for files not recently used. During the time that Integrity Server 100 is standing in for failed server 202, it continues to capture and manage protection copies of the files of other servers 102. When the failed server 202 is recovered and brought back on line, files are synchronized so that no data are lost.

Detailed Description Text - DETX (13):

Referring again to FIG. 1, Integrity Server 100 has a disk 120, a tape auto-loader, and runs Novell NetWare version 4.10 or later, a client/server communications system (TIRPC), and a file transport system (Novell SMS). An example tape auto-loader 110 is an HP 1553c, that holds six 8 GB tapes.

Detailed Description Text - DETX (14):

Each protected server 102 runs Novell NetWare, version 3.11 or later, TIRPC, Novell SMS components appropriate to the NetWare version, and runs an agent program for copying the modified files.

Detailed Description Text - DETX (17):

Referring again to FIG. 1, in protection mode, Integrity Server 100 manages its data store to meet several objectives. The most actively used data are kept in the disk cache 120, so that when the Integrity Server is called on to stand in for a server 102, the most active files are available from disk cache 120. All current files from all protected servers 102 are kept on tape, available for automatic retrieval to the disk cache for use during stand-in, or for conventional file restoration. A set of tapes is created and maintained for off-site storage to permit recovery of the protected servers and the Integrity Server itself if both are destroyed or inaccessible. All files stored on tape are stored twice before the disk copy is removed, once on active tape 150 and once on offsite tape 164.

Detailed Description Text - DETX (18):

A continuously protected system usually has the following tapes in its autoloader(s): a current active tape 150, the rest of the filled active tapes

151-153 of the active set, possibly an active tape that the Integrity Server has asked the System Manager to dismount and file in legacy storage, one current offsite tape 164, possibly a recently-filled off-site tape, possibly a cleaning tape, and several blank (or overwritable) tapes.

Detailed Description Text - DETX (19):

The server agents and Integrity Server 100 maintain continuous communication, with the agents polling the Integrity Server for instructions, and copying files. Based on a collection of rules and schedules selected by the system manager, agents perform tasks on a continuous, scheduled, or demand basis. Each agent continuously scans the directories of its server looking for new or changed files, detected, for example, using the file's NetWare archive bit or its last modified date/time stamp. Similarly, newly-created files are detected and copied to the Integrity Server. In normal operation, a single scan of the directories of a server takes on the order of fifteen minutes. If a file changes several times within this protection interval, only the most recent change will be detected and copied to the Integrity Server. A changed file need not be closed to be copied to the Integrity Server, but it must be sharable. Changes made to non-sharable files are protected only when the file is closed.

Detailed Description Text - DETX (20):

In one embodiment, the protected server's protection agent registers with the NetWare file system's File System Monitor feature. This registration requests that the agent be notified when a client requests a file open operation, prior to the file system's execution of the open operation. When a Protected Server's protection agent opens a file, the file is opened in an exclusive mode so that no other process can alter the file before an integral snapshot is sent to the Integrity Server. Further, the agent maintains a list of those files held open by the agent, rather than, e.g., on behalf of a client. When a client opens a file, the protection agent is notified by the File System Monitor and consults the list to determine if the agent currently has the file open for snapshotting to the Integrity Server. While the agent has the file open, the client process is blocked (that is, the client is held suspended) until the agent completes its copy operation. When the agent completes its snapshot, the client is allowed to proceed. Similarly, if the agent does not currently have the file open, a client request to open a file proceeds normally.

Detailed Description Text - DETX (21):

When an agent process of one of the file servers detects a file update on a protected server 102, the agent copies the file new version of the changed file and related system data to the Integrity Server's disk cache 120. (As a special case, when protection is first activated, the agent walks the server's directory tree and copies all files designated for protection to the Integrity Server.) The Integrity Server queues the copied file in the active queue 142 and then off-site queue 160 for copying to the active tape 150 and off-site tape 164, respectively. Some files may be scheduled for automatic periodic copying from server 102 to Integrity Server 100, rather than continuous protection.

Detailed Description Text - DETX (22):

The population of files in the disk cache 120 is managed to meet several desired criteria. The inviolable criterion is that the most-recent version of a file sampled by the server's agent process always be available either in disk cache 120 or on one of the tapes 150-153,164 of the autoloader. Secondary

criteria include reducing the number of versions retained in the system, and maintaining versions of the most actively used files on the disk cache so that they will be rapidly ready for stand-in operation.

Detailed Description Text - DETX (23):

A given file version will be retained in disk cache 120 for at least the time that it takes for the version to work its way through active queue 142 to active tape 150, and through offsite queue 160 for copying to current off-site tape 164. Once a file version has been copied to both the active and off-site tapes, it may be kept on disk 120 simply to provide the quickest possible access in case of failure of the file's protected server. The version may be retained until the disk cache 120 approaches being full, and then the least active file versions that have already been saved to both tapes are purged.

Detailed Description Text - DETX (24):

Redundant versions of files are not required to be stored in cache 120. Thus, when a new version of a protected file is completely copied to disk cache 120, any previous version stored in cache 120 can be erased (unless, for instance, that version is still busy, for instance because it is currently being copied to tape). When a new version displaces a prior version, the new history package is left at the tail of the active queue so that the file will be retained in disk cache 120 for the maximum amount of time. As files are dequeued from active queue 142 for copying to active tape 150, the most-recent version of the file already in the disk cache is written to tape, and all older versions are removed from the queue.

Detailed Description Text - DETX (25):

The active tape set 150-153 and the data stored thereon is actively managed by software running on Integrity Server 100, to keep the most recent file versions readily available on a small number of tapes. Data are reclaimed from the oldest active tape 153 and compacted so that the oldest active tape can be removed from the autoloader for storage as a legacy tape 168. Compaction is triggered when the density of the data (the proportion of the versions on the active tape that have not been superseded by more-recent versions, eg. in the disk cache or later in the active tape set), averaged across all active tapes 150-153 currently in the autoloader, falls below a predetermined threshold (e.g. 70%), or when the number of available blank tapes in autoloader 110 falls below a threshold (e.g., 2). In the compaction process, the file versions on oldest active tape 153 that are up to date with the copy on the protected server, and thus which have no later versions in either disk cache 120 or on a newer active tape 150-152, are reclaimed by copying them from oldest active tape 153 to the disk cache 120 (unless the file version has been retained in disk cache 120). From disk cache 120, the version is re-queued for writing to a new active tape 150 and off-site tape 164, in the same manner as described above for newly-modified files. This re-queuing ensures that even read-active (and seldom-modified) data appear frequently enough on active tapes 150 and off-site tapes 165 to complete a restorable set of all protected files. Since all data on oldest active tape 153 are now either obsolete or replicated elsewhere 120,150-152 on Integrity Server 100, the tape 153 itself may now be removed from the autoloader for retention as a legacy tape 168.

Detailed Description Text - DETX (26):

The compaction process ensures that every protected file has an up-to-date copy accessible from the active tape set. Once the active tape set has been compacted, i.e., current files have been copied from the oldest active tape 153 to the newest active tape 150 and an off-site tape 164, the oldest active tape

is designated a legacy tape 168, and is ready to be removed from the autoloader. Its slot can be filled with a blank or expired tape.

Detailed Description Text - DETX (27):

The process of reclamation and compaction does not change the contents of the oldest active tape 153. All of its files remain intact and continue to be listed in the Integrity Server's catalog. A legacy tape and its files are kept available for restoration requests, according to a retention policy specified by the system manager. Legacy tapes are stored, usually on-site, under a user-defined rotation policy. When a legacy tape expires, the Integrity Server software removes all references to the tape's files from the catalog. The legacy tape can now be recycled as a blank tape for reuse as an active or off-site tape. The Integrity Server maintains a history of the number of times each tape is reused, and notifies the system manager when a particular tape should be discarded.

Detailed Description Text - DETX (28):

Note that the process of reclaiming data from the oldest active tape 153 to disk cache 120 and then compacting older, non-superseded versions to active tape 150 allows the Integrity Server 100 to maintain an up-to-date version of a large number of files, exploiting the low cost of tape storage, while keeping bounded the number of tapes required for such storage, without requiring periodic recopying of the files from protected servers 102. The current set of active tapes should remain in the autoloader at all times so that they can be used to reconstruct the stored files of a failed server, though the members of the active tape set change over time.

Detailed Description Text - DETX (29):

By ensuring that every protected file is copied to offsite tape 164 with a given minimum frequency (expressed either in time, or in length of tape between instances of the protected file), the process also ensures that the offsite tapes 165 can be compacted, without physically accessing the offsite tape volumes.

Detailed Description Text - DETX (30):

In an alternate tape management strategy, after reclaiming the still-current file versions from oldest active tape 153, this tape is immediately recycled as the new active tape 150. This forgoes the benefit of the legacy tapes' maintenance of recent file versions, but reduces human intervention required to load and unload tapes.

Detailed Description Text - DETX (31):

Writing files from the off-site queue 160 to off-site tape 164 is usually done at low priority, and the same version culling described for active queue 142 is applied to off-site queue 160. The relatively long delay before file versions are written to off-site tape 164 results in fewer versions of a rapidly-changing file being written to the off-site tape 164, because more of the queued versions are superseded by newer versions.

Detailed Description Text - DETX (32):

Whether it has been updated or not, at least one version of every protected file is written to an off-site tape with a maximum number of sequential off-site tapes between copies. This ensures that every file appears on at least every $n.\text{sup.th}$ tape (for some small n), and ensures that any sequence of

n consecutive off-site tapes contains at least one copy of every protected file, and thus that the sequence can serve the function of a traditional backup tape set, providing a recovery of the server's files as they stood at any given time.

Detailed Description Text - DETX (34):

Even though off-site tapes are individually removed from the autoloader and individually sent off-site for storage, successive tapes together form a "recovery set" that can be used to restore the state of the Integrity Server in case of disaster. The circularity of the tape compaction process ensures that at least one version of every file is written to an off-site tape with a maximum number of off-site tapes intervening between copies of the file, and thus that a small number of consecutive off-site tapes will contain at least one version of every protected file. To simplify the process of recovery, the set of off-site tapes that must be loaded to the Integrity Server to fully recover all protected data is dynamically calculated by the Integrity Server at each active tape compaction, and the tape ID numbers of the recovery set ending with each off-site tape can be printed on the label generated as the off-site tape is removed from the autoloader. When a recovery is required, the system manager simply pulls the latest off-site tape from the vault, and also the tapes listed on that tape's label, to obtain a set of off-site tapes for a complete recovery set.

Detailed Description Text - DETX (35):

Many tape read errors can be recovered from with no loss of data, because many file versions are redundantly stored on the tapes (e.g., a failure on an active tape may be recoverable from a copy stored on an off-site tape).

Detailed Description Text - DETX (37):

Expired off-site tapes cannot be used to satisfy file restoration requests, because the history packages for the tape will have been purged from the catalog. But these tapes may still be used for Integrity Server recovery, as long as a full recovery set is available and all tapes in the set can be read without error.

Detailed Description Text - DETX (38):

The history packages are maintained on disk 120, rather than in the RAM of the Integrity Server, so that they will survive a reboot of the Integrity Server. The history packages are linked in two ways. Active queue 142 and off-site queue 160 are maintained as lists of history packages, and the history packages are also maintained in a tree structure isomorphic to the directory tree structure of the protected file systems. Using the tree structure, a history package can be accessed quickly if the file version needs to be retrieved from either the active tape set 150-153 or from an off-site tape, either because Integrity Server 100 has been called to stand in for a failed server, or because a user has requested a restore of a corrupted file.

Detailed Description Text - DETX (39):

File versions that have been copied to both active tape 150 and off-site tape 164 can be erased from disk cache 120. In one strategy, files are only purged from disk cache 120 when the disk approaches full. Files are purged in least-recently accessed order. It may also be desirable to keep a most-recent version of certain frequently-read (but infrequently-written) files in disk cache 120, to provide the fastest-possible access to these files in case of server failure.

Detailed Description Text - DETX (40):

Depending on which tape (an active tape 150 or an off-site tape 164) is loaded into the autoloader's read/write station and the current processing load of the Integrity Server, a given file version may take anywhere from a few minutes to hours to be stored to tape. The maximum time bound is controlled by the System Manager. Typically a file version is stored to active tape 150 as quickly as possible, and queued for the off-site tape at a lower priority.

Detailed Description Text - DETX (41):

Verification of tape writes may be enabled by the System Manager Interface. When tape write verification is enabled, each queue is fully written to tape, and then the data on the tape are verified against the data in disk cache 120. Files are not requeued from the active tape queue 142 to the off-site queue 160 until the complete active tape 150 is written and verified.

Detailed Description Text - DETX (43):

In some embodiments, a System Manager can request that a specified file be protected within a specific time window, such as when there is no update in progress or when the file can be closed for protection purposes.

Detailed Description Text - DETX (44):

Referring to FIGS. 3a and 3b, a catalog records where in the Integrity Server (e.g, on disk 120, active tapes 150-153, legacy tapes 168, or off-site tapes 164-165) a given file version is to be found. It contains detailed information about the current version of every file, such as its full filename, timestamp information, file size, security information, etc. Catalog entries are created during protection mode as each file version is copied from the protected server to the Integrity Server. Catalog entries are altered in form and storage location as the file version moves from disk cache 120 to tape and back. The catalog is used as a directory to the current tapes 150-153, legacy tapes, and off-site tapes 164 when a user requests restoration of or access to a given file version.

Detailed Description Text - DETX (45):

FIGS. 3a and 3b show two data structures that make up the catalog. The catalog has entries corresponding to each leaf file, each directory, each volume, and each protected server, connected in trees corresponding to the directory trees of the protected servers. Each leaf file is represented as a single "file package" data structure 310 holding the stable properties of the file. Each file package 310 has associated with it one or more "history package" data structures 312, each corresponding to a version of the file. A file package 310 records the file's creation, last access, last archive date/time, and protection rights. A history package 312 records the location in the Integrity Server's file system, the location 316 on tape of the file version, the date/time that this version was created, its size, and a data checksum of the file contents. Similarly, each directory and volume have a corresponding data structure. As a version moves within the Integrity Server (for instance, from disk cache 120 to tape 150-153), the location mark 316 in the history package is updated to track the files and versions.

Detailed Description Text - DETX (46):

Other events in the "life" of a file are recorded in the catalog by history packages associated with the file's file package. Delete packages record that

the file was deleted from the protected server at a given time (even though one or more back versions of the file are retained by the Integrity Server).

Claims Text - CLTX (1):

1. A method for managing copies of a protected set of files on a bounded number of sequential-access volumes, the method being executed by computer and comprising:

Claims Text - CLTX (3):

(b) when an external process independent of the sequential-access volumes alters the contents of one of the protected files to produce a new current version of the protected file, snapshotting the new current version of the altered protected file at the end of the current volume;

Claims Text - CLTX (5):

(d) maintaining the population of an active set of said sequential-access volumes at or below said bounded number, said active set being the minimum set of the most-recently-current of said volumes that together contain at least one version of each of said protected files, by:

Claims Text - CLTX (7):

copying from the compaction volume to the current volume those versions of file versions stored on the compaction volume not having a more recent version stored on the active set; and

Claims Text - CLTX (10):

when the contents of one of the protected files is altered to produce a new current version of the protected file, snapshotting the new current version of the altered protected file to a direct access storage cache; and

Claims Text - CLTX (11):

queueing in a write queue the cache copy of the file for later writing to said current volume.

Claims Text - CLTX (13):

when dequeuing a file from said write queue for writing to said current volume, reviewing said write queue for a later version of the dequeued file, and suppressing writing to said current volume of any version other than the latest queued version of said dequeued file.

Claims Text - CLTX (15):

copying a file version from the compaction volume to said storage cache; and

Claims Text - CLTX (16):

queueing the file from the compaction volume in a write queue as the cache copy of the file for later writing to said current volume.

Claims Text - CLTX (18):

when dequeuing a file from said write queue for writing to said current

volume, reviewing said write queue for a later version of the dequeued file, and suppressing writing of any version other than the latest queued version of said dequeued file.

Claims Text - CLTX (20):

recently-compacted volumes are maintained as a legacy set of volumes containing additional copies of current versions and non-current versions of protected files, and

Claims Text - CLTX (21):

storage records corresponding to the file versions stored on the legacy set are retained allowing prompt retrieval of those copies as requested by the external process.

Claims Text - CLTX (23):

concurrently with steps (b)-(c), enqueueing versions of said altered files for writing to an archival set of volumes distinct from said active and legacy volumes;

Claims Text - CLTX (24):

wherein all volumes of the active volume set are resident in an auto-loader, and all file versions of said active volume set can be retrieved with a relatively small latency; and

Claims Text - CLTX (27):

after a file version is written to the current volume, queuing said file version to be written to the current archival volume in an archival queue.

Claims Text - CLTX (29):

retaining a file version in said archival queue for a time; and

Claims Text - CLTX (30):

when dequeuing a file from said archival queue for writing to said current archival volume, reviewing said archival queue for a later version of the dequeued file, and suppressing writing to said current archival volume of any version other than the latest queued version of said dequeued file.

Claims Text - CLTX (32):

maintaining at a small number a population of said archival volumes preceding each said archival volume that taken together form a recovery set, a recovery set being a sequence of consecutive ones of said archival volumes that collectively contain at least one version of every file of the protected set, by copying to said current archival volume those file versions copied from the compaction volume to the current volume.

Claims Text - CLTX (38):

the altering of the file servers' protected files includes creation of a file by the external process.

Claims Text - CLTX (39):

14. A method for protecting a protected set of files of varying size and stored on direct-access mass storage devices of a plurality of file server nodes of a network of computers, the method comprising:

Claims Text - CLTX (40):

at a rate similar to the rate at which said files are altered by an external process, snapshotting recently-altered protected ones of said files from said direct-access mass storage devices to an archive storage cache, a new snapshot of a given file in said storage cache displacing any older snapshot of said given file in existence in said storage cache;

Claims Text - CLTX (43):

periodically verifying the contents of the protected files against the contents of the versions stored on said removable mass storage media.

Claims Text - CLTX (45):

avoiding copying contents of verified files over the network, or reading the removable media, during the verifying, by comparing a summary value of the content of a protected file with a summary value of the content of the stored snapshot.

Claims Text - CLTX (48):

18. The method of claim 17 wherein an active set, being a minimum set of most-recently-written volumes of said media that together contain at least one version of each of said protected files, is maintained at a bounded number of volumes, by further steps comprising:

Claims Text - CLTX (50):

(b) when the contents of one of the files is altered to produce a new current version of the file, snapshotting the new current version of the altered file to the current volume;

Claims Text - CLTX (54):

copying from the compaction volume to the current volume those versions of file versions stored on the compaction volume not having a more recent version stored on the active set; and

Claims Text - CLTX (57):

from among the volumes previously compacted, periodically expiring some of said volumes, leaving short sequences of consecutively-generated volumes that taken together store at least one copy of every file of said protected set.

Claims Text - CLTX (61):

maintaining at a small number of volumes a population of said volumes of said off-site media preceding each said off-site volume that taken together form a recovery set, a recovery set being a sequence of consecutive ones of said off-site volumes that collectively contain at least one version of every file of the protected set, by periodically refreshing said storage cache with copies of protected files not recently altered, and copying said refreshed files from said storage cache to said removable mass storage media.

Claims Text - CLTX (63):

periodically expiring media from said off-site archive, leaving short sequences of consecutively-generated volumes of said off-site archive that taken together store at least one copy of every file of said protected set.

Claims Text - CLTX (65):

the altering of the file servers' files by the external process includes creation of new files in the protected set by the external process, and said newly-created ones of said protected files are snapshotted to said archive storage cache.

Claims Text - CLTX (67):

recording that a protection process holds the file open during said snapshotting; and

Claims Text - CLTX (68):

when said external process requests access to a file, consulting said recording to determine whether the file is currently held open by said protection process, and:

Claims Text - CLTX (69):

if the file is currently held open by said protection process, blocking said external process until said protection process completes snapshotting of the file, and

Claims Text - CLTX (70):

if the file is not currently held open by said protection process, or when the file is released by the protection process, proceeding to open the file in accord with the file open protocol of said protected computer.

Claims Text - CLTX (72):

traversing a file system of a protected computer by a protection process, snapshotting files of said file system to removable storage media, and

Claims Text - CLTX (73):

as each file is opened for said snapshotting, recording that the protection process currently holds the file open, and

Claims Text - CLTX (74):

as the protection process completes said snapshotting, recording that the protection process has released the file; and

Claims Text - CLTX (75):

when a client process requests access to a file, consulting said recording to determine whether the file is currently open by said protection process, and

Claims Text - CLTX (76):

if the file is currently held open by said protection process, blocking said client process until said protection process releases the file, and

Claims Text - CLTX (77):

if the file is not currently held open by said protection process, or when the protection process completes said snapshotting, proceeding to open the file in accord with the file open protocol of said protected computer.

Claims Text - CLTX (82):

said sequential-access volumes are removable mass storage media in a form suitable for an archive for off-site storage, said snapshotting being carried out at a rate similar to the rate at which said files are altered by an external process, a new snapshot of a given file in said storage cache displacing any older snapshot of said given file in existence in said storage cache;

Claims Text - CLTX (85):

maintaining at a small number of volumes a population of said volumes of said off-site media preceding each said off-site volume that taken together form a recovery set, a recovery set being a sequence of consecutive ones of said off-site volumes that collectively contain at least one version of every file of the protected set, by periodically refreshing said storage cache with copies of protected files not recently altered, and copying said refreshed files from said storage cache to said removable mass storage media.

Claims Text - CLTX (87):

periodically expiring media from said off-site archive, leaving short sequences of consecutively-generated volumes of said off-site archive that taken together store at least one copy of every file of said protected set.

Claims Text - CLTX (89):

the altering of the protected set of files by the external process includes creation of new files in the protected set by the external process, and said newly-created ones of said protected files are snapshotted to said storage cache.

Claims Text - CLTX (91):

recording that a protection process holds the file open during said snapshotting; and

Claims Text - CLTX (92):

when said external process requests access to a file, consulting said recording to determine whether the file is currently held open by said protection process, and:

Claims Text - CLTX (93):

if the file is currently held open by said protection process, blocking said external process until said protection process completes snapshotting of the file, and

Claims Text - CLTX (94):

if the file is not currently held open by said protection process, or when the file is released by the protection process, proceeding to open the file in accord with the file open protocol of said protected computer.

Other Reference Publication - OREF (5):

White Paper, "St. Bernard Software Open File Manager", Mar. 7, 1995.